

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Bavaria et al.	§
Serial No.: 10/631,056	§ Group Art Unit: 2144
Filed: July 31, 2003	§ Examiner: Ibrahim, Mohamed
For: Method and Apparatus for	§ Confirmation No.: 3503
Performing Device Configuration	§
Rediscovery	§

35525

PATENT TRADEMARK OFFICE
CUSTOMER NUMBER

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF (37 C.F.R. 41.37)

This brief is in furtherance of the Notice of Appeal, filed in this case on January 9, 2009.

A fee of \$540.00 is required for filing an Appeal Brief. Please charge this fee to IBM Corporation Deposit Account No. 09-0447. No additional fees are believed to be necessary. If, however, any additional fees are required, I authorize the Commissioner to charge these fees which may be required to IBM Corporation Deposit Account No. 09-0447.

REAL PARTY IN INTEREST

The real party in interest in this appeal is the following party: International Business Machines Corporation of Armonk, New York.

RELATED APPEALS AND INTERFERENCES

This appeal has no related proceedings or interferences.

STATUS OF CLAIMS

A. TOTAL NUMBER OF CLAIMS IN APPLICATION

The claims in the application are: 1-27

B. STATUS OF ALL THE CLAIMS IN APPLICATION

Claims canceled: 6-20

Claims withdrawn from consideration but not canceled: None

Claims pending: 1-5 and 21-27

Claims allowed: None

Claims rejected: 1-5 and 21-27

Claims objected to: None

C. CLAIMS ON APPEAL

The claims on appeal are: 1-5 and 21-27

STATUS OF AMENDMENTS

No amendment after final rejection was filed for this case.

SUMMARY OF CLAIMED SUBJECT MATTER

The present invention is generally directed to techniques to efficiently determine and process device configuration information in a data processing system. Such efficiency is achieved both in time (reduced time for determining configuration) and in space (reduced storage requirements for maintaining configuration code). A two-pronged approach is used – an initial discovery phase and a rediscovery phase - and due to the unique architecture, *the same configuration code is operable in performing both of these different phases of device configuration* – with such code re-use capability being made possible by unique memory management operations where configuration data is conditionally moved from one memory location to another memory location under certain conditions (a match) and configuration data is read from a device to the another memory location under certain other conditions (no match). In the initial discovery phase, device configuration information is read from the devices (initial device ID) and maintained in a memory region (stored device ID) for subsequent use. In the rediscovery phase, a table is used to access unique device identification information from the devices (current device ID) and compared with unique device identification information maintained in the memory as per the initial discovery phase (stored device ID). *If a match exists between the current device ID and the stored device ID, the stored device ID is moved to another memory/region (moved device ID) such that the previously read device ID (stored device ID) is no longer present in the first memory region* (Specification page 13, first two paragraphs). This moving of the device ID to another memory region facilitates use of the same configuration code for both phases (discovery and rediscovery), as one phase (discovery phase) uses one memory region (buffer memory region 1 of Figure 4) to facilitate the configuration process, whereas the other phase (rediscovery phase) uses the other memory region (temporary memory region 3 of Figure 4) to temporarily store configuration data such that the first memory region can be re-used during completion of the rediscovery phase to facilitate the configuration process (Specification page 14, first paragraph). The teachings of the cited reference do not contemplate such memory usage and associated data relocation, which advantageously allows for re-using the same configuration code for both phases (discovery and rediscovery), thereby reducing the requisite amount of firmware code, and its associated (memory/storage) space (Specification page 14, lines 8-12; page 20, lines 6-13 and lines 27-30), in addition to reducing the overall time required to determine configuration information for the devices in the data processing system.

A. CLAIM 1 - INDEPENDENT

The subject matter of claim 1 is directed to a method in a data processing system for identifying device configurations (Specification page 12, lines 1-4). The method comprises identifying unique identification information for a set of devices in the data processing system to form identified unique identification information (Specification page 13, lines 5-8; page 16, lines 23-24). The method comprises comparing the identified unique identification information with previously identified unique identification information (Specification page 13, lines 8-10; page 16, lines 24-25; page 18, lines 3-5; Figure 5, element 500). The method also comprises moving configuration data to a memory for devices in the set of devices in which a match exists between the identified unique identification information and the previously identified unique identification information for devices (Specification page 13, lines 10-13; page 16, lines 25-29; page 18, lines 5-7; Figure 5, element 502; Figure 6, elements 608 and 610). The method comprises obtaining configuration information from a device in which configuration information is absent in the memory after configuration data has been moved to the memory for the devices to form a current set of configuration data for the set of devices (Specification page 13, line 30 – page 14, line 3; page 17, lines 19-21; page 18, lines 7-9; Figure 5, element 504), wherein the previously identified unique identification information is accessed using a table associated with the configuration data for the set of devices (Specification page 12, lines 25-26; page 15, lines 25-26; Figure 4, element 410), wherein the table comprises (i) an index used to locate particular configuration data for a particular device (Specification page 15, lines 27-29), (ii) information used to address the particular device (Specification page 15, lines 29-30), and (iii) an offset to a memory location within the particular device at which particular unique identifier information for the particular device is stored (Specification page 13, lines 1-4; page 16, lines 1-3).

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

The grounds of rejection to review on appeal are as follows:

A. GROUND OF REJECTION 1

The rejection of Claims 1 and 3-5 under 35 U.S.C. § 102 as being anticipated by Kartoz (U.S. Patent No. 7,024,547);

B. GROUND OF REJECTION 2

The rejection of Claim 2 under 35 U.S.C. § 103 as being unpatentable over Kartoz (U.S. Patent No. 7,024,547) in view of Zintel (U.S. Patent No. 6,779,004);

C. GROUND OF REJECTION 3

The rejection of Claim 2 under 35 U.S.C. § 103 as being unpatentable over Kartoz (U.S. Patent No. 7,024,547) in view of Zintel (U.S. Patent No. 6,779,004) and further in view of Krejsa (U.S. Publication No. 2004/0107329); and

D. GROUND OF REJECTION 4

The rejection of Claims 21-27 under 35 U.S.C. § 103 as being unpatentable over Kartoz (U.S. Patent No. 7,024,547) in view of Garney (U.S. Patent No. 5,854,905).

ARGUMENT

A. GROUND OF REJECTION 1 (Claims 1 and 3-5)

Claims 1 and 3-5 stand rejected under 35 U.S.C. § 102 as being anticipated by Kartoz (U.S. Patent No. 7,024,547), hereinafter “Kartoz”.

1. *Claims 1 and 3-5*

For a prior art reference to anticipate in terms of 35 U.S.C. 102, every element of the claimed invention must be identically shown in a single reference. *In re Bond*, 910 F.2d 831, 15 USPQ2d 1566 (Fed. Cir. 1990). Section 102 embodies the concept of novelty—if a device or process has been previously invented (and disclosed to the public), then it is not new, and therefore the claimed invention is "anticipated" by the prior invention. . . . Because the hallmark of anticipation is prior invention, the prior art reference—in order to anticipate under 35 U.S.C. § 102—must not only disclose all elements of the claim within the four corners of the document, but must also disclose those elements "arranged as in the claim." *Net MoneyIn v. Verisign* (No. 07-1565, Fed. Cir., October 2008). Appellants will now show that every element recited in Claim 1, including (i) a conditional moving operation, and (ii) details of device configuration data that is maintained in a table, is not identically shown - nor arranged as in the claim - in the cited Kartoz reference, and thus Claim 1 has been erroneously rejected under 35 U.S.C. § 102.

With respect to Claim 1, such claim recites “moving configuration data to a memory for devices in the set of devices in which a match exists between the identified unique identification information and the previously identified unique identification information for devices”. As can be seen, *when a match exists* between the identified unique identification information and the previously identified unique identification information for a device, *configuration data for such device is moved to a memory*. In rejecting this aspect of Claim 1, the Examiner states that this moving of configuration data is taught by the cited Kartoz reference since such reference teaches at col. 4, line 64 – col. 5, line 4:

“if the reference identification data of the device **match**, the system *uses* the reference initialization data *to initialize the device*”(emphasis added by Appellants)

Appellants respectfully submit that this *use* of reference initialization data to initialize a device, as described by the cited reference, does not teach any type of ‘moving’ operation¹ where *configuration data for a device is conditionally moved to a memory* if a match exists between the identified unique identification information and the previously identified unique identification information for devices, as per the features of Claim 1.² Importantly, **Kartoz is keen on NOT performing additional operations if a match occurs (Kartoz col. 5, lines 1-4)**, instead merely *using* reference data previously acquired during a previous boot procedure (Kartoz col. 4, lines 5-8). Thus, in addition to not being anticipated, it is further urged that Claim 1 is not obvious in view of the cited reference as a person of ordinary skill in the art would not have been motivated to conditionally perform data move operations in response to a match that Kartoz disdains.

In addition, per Claim 1 this configuration data is moved to a different item (memory) than the item for which the configuration data pertains to (matching device), due to the fact that ‘device’ and ‘memory’ are both used in Claim 1, and thus configuration information for one item (the matching device) is moved to another item (the memory). In contrast, this cited Kartoz passage describes *using* reference initialization data for a device to *initialize the same device*. Thus, not only is there no teaching of *moving* configuration data upon occurrence of a match, there is also no teaching of moving configuration data associated with one item (a matching device) to another item (a memory). Therefore, there are at least two claimed features that are not identically shown in the cited reference. First, the cited reference does not teach *moving* of configuration data upon occurrence of a device match. Second, the cited reference does not teach moving configuration data to *a different item* (a memory) that was not the subject of the ‘match’

¹ **Move:** to change from one place or position to another (Source: <http://www.dictionary.com>)

² During examination, the claims must be interpreted as broadly as their terms reasonably allow. *In re American Academy of Science Tech Center*, 367 F.3d 1359, 1369, 70 USPQ2d 1827, 1834 (Fed. Cir. 2004). This means that the words of the claim must be given their **plain meaning** unless the plain meaning is inconsistent with the specification. *In re Zletz*, 893 F.2d 319, 321, 13 USPQ2d 1320, 1322 (Fed. Cir. 1989) (discussed below); *Chef America, Inc. v. Lamb-Weston, Inc.*, 358 F.3d 1371, 1372, 69 USPQ2d 1857 (Fed. Cir. 2004) (Ordinary, simple English words whose meaning is clear and unquestionable, absent any indication that their use in a particular context changes their meaning, are **construed to mean exactly what they say**. MPEP 2111.01(I) (emphasis added by Appellants). The Examiner is not interpreting the claim terms in accordance with their normal, plain meaning.

determination (a device), but instead teaches that the same item (memory device) is initialized when this same item (memory device) is ‘matched’. Thus, both (1) the ‘operation’ performed (‘using’ per the cited reference versus ‘moving’ per Claim 1) and (2) the item/thing that the operation is performed upon (the same device as what was ‘matched’ per the teachings of the cited reference, versus a different item/thing (a memory) from what was ‘matched’ per Claim 1) are different between what is recited in Claim 1 and what is taught by the cited reference. Accordingly, as *every element* is not *identically shown and identically arranged* in a single reference, it is urged that Claim 1 is not anticipated by the cited reference.

As described above in the introductory Summary of Claimed Subject Matter section, this moving of configuration data (such movement being from area 1 to area 3 in the preferred embodiment, as described at Specification page 16, lines 25-29 and depicted in Figure 4) allows for a temporary relocation of the configuration data to a temporary storage area so the source area can be re-used during the rediscovery process (Specification page 16, lines 6-8), which advantageously facilitates use of the configuration code for both the discovery phase as well as the rediscovery phase (Specification page 14, lines 8-12; page 20, lines 6-13 and lines 27-30), in that the first memory region (area 1) is effectively re-initialized at the end of this device match processing due in-part to such moving (Specification page 17, lines 7-10), thereby allowing the first memory region to be used as if it were in an initialized state during the remainder of the rediscovery phase (Specification page 17, lines 11-25) – and thus allowing for re-use of the discovery code during both discovery phases. The cited reference does not teach (or otherwise suggest) these claimed features or their resulting advantages. Therefore, it is urged that Claim 1 has been erroneously rejected under 35 U.S.C. § 102, as there are at least two claimed features that are not identically shown in the cited reference.

Further with respect to Claim 1, such claim recites “wherein the previously identified unique identification information is accessed using a table associated with the configuration data for the set of devices, wherein the table comprises (i) an index used to locate particular configuration data for a particular device, (ii) information used to address the particular device, and (iii) an offset to a memory location within the particular device at which particular unique identifier information for the particular device is stored”. As can be seen, a table associated with configuration data for the set of devices is used to access the *previously identified unique identification information*, with this table comprising (1) an index used to locate particular

configuration data for a particular device, (2) information used to address the particular device, and (3) an offset to a memory location within the particular device at which particular unique identifier information for the particular device is stored. It is urged that the table that is alleged to be taught by the cited Kartozy reference does not include each of these three (3) explicitly enumerated items - and these three (3) enumerated items advantageously facilitate access to particular parameters within the devices such as configuration data and unique identifier information for such devices (Specification page 15, bottom – page 16, top). In contrast, the alleged Kartozy table (per the cited teachings of Kartozy at col. 4, lines 44-55) includes parameters that are used to *initialize* a device (Kartozy col. 4, lines 49-51) and reference identification data that is used to determine whether the memory configuration has changed (Kartozy col. 4, lines 52-55). None of this data pertains to (i) an index used to locate particular configuration data for a particular device, (ii) information used to address the particular device, and (iii) an offset to a memory location within the particular device at which particular unique identifier information for the particular device is stored.

Importantly, *the Examiner themselves admits that the table information is different between what is claimed and what is taught by the cited Kartozy reference*³, and thus the Examiner in essence admits that every element recited in Claim 1 is not identically shown and identically arranged in a single reference – and therefore Kartozy does not satisfy the all-elements-rule “arranged as in the claim” that must be complied with to make a proper 35 U.S.C. § 102 rejection.⁴

Thus, it is further shown that Claim 1 is not anticipated by the cited reference, and has thus been erroneously rejected 35 U.S.C. § 102, as there are additional claimed elements that are not identically shown and not identically arranged in a single reference.

³ See, e.g., page 8, lines 1-3 of the Final Office Action dated July 11, 2008.

⁴ Nor is the information in the table a mere intended use, as alleged by the Examiner, as such claim recites “wherein the previously identified unique identification information is accessed using a table associated with the configuration data for the set of devices, wherein the table comprises (i) an index used to locate particular configuration data for a particular device, (ii) information used to address the particular device”. Also see Claim 21 for further usage of such table.

B. GROUND OF REJECTION 2 (Claim 2)

Claim 2 stands rejected under 35 U.S.C. § 103 as being unpatentable over Kartozi in view of Zintel (U.S. Patent No. 6,779,004), hereinafter “Zintel”.

1. Claim 2

Appellants initially urge error in the rejection of Claim 2 for similar reasons to those given above with respect to Claim 1, as the newly cited Zintel reference does not overcome the teaching deficiencies identified hereinabove.

Further with respect to Claim 2, such claim recites “wherein the memory is a temporary random access memory comprising an area for maintaining the previously identified unique identification information and another area for maintaining the moved configuration data while the obtaining configuration information is being performed”. As can be seen, the memory has **two distinct areas – (1) an area for maintaining the previously identified unique identification information, and (2) another area for maintaining the moved configuration data while the obtaining configuration information is being performed**. It is urged that none of the cited references teach or suggest a memory having these two distinct areas – one area for maintaining previously identified information and another area for maintaining the moved configuration data while the obtaining configuration information is being performed. In rejecting Claim 2, the Examiner merely alleges that Zintel discloses ‘system for auto-configuring of peripherals that stores unique device identifier *in random access memory*’. The Examiner fails to address or allege any teaching or suggestion that the cited references describe the particular memory configuration recited in Claim 2 – and in fact the Examiner explicitly admits that these cited references do NOT teach these claimed features⁵ - and therefore Claim 2 has been erroneously rejected as a proper prima facie showing of obviousness has not been established, as such claim does not merely recite a RAM device having unique identifiers stored therein, as alleged by the Examiner in rejecting such claim.⁶

⁵ The Examiner themselves admits that this Kartozi/Zintel combination does not teach the different area features of Claim 2 (see, e.g., page 5, paragraph 5 of the Final Office Action dated July 11, 2008).

⁶ In rejecting claims under 35 U.S.C. Section 103, the examiner bears the initial burden of presenting a prima facie case of obviousness. *In re Oetiker*, 977 F.2d 1443, 1445, 24 USPQ2d

C. GROUND OF REJECTION 3 (Claim 2)

Claim 2 stands rejected under 35 U.S.C. § 103 as being unpatentable over Kartozi in view of Zintel and further in view of Krejsa (U.S. Publication No. 2004/0107329), hereinafter “Krejsa”.

1. Claim 2

Appellants initially urge error in the rejection of Claim 2 for similar reasons to those given above with respect to Claim 1, as the newly cited Zintel and Krejsa references do not overcome the teaching deficiencies identified hereinabove.

Further with respect to Claim 2, such claim recites “wherein the memory is a temporary random access memory comprising an area for maintaining the previously identified unique identification information and another area for maintaining the moved configuration data while the obtaining configuration information is being performed”. As can be seen, the memory has two distinct areas – (1) an area for maintaining the previously identified unique identification information, and (2) another area for maintaining the moved configuration data while the obtaining configuration information is being performed. It is urged that none of the cited references teach or suggest a memory having these two distinct areas – one area for maintaining previously identified information and another area for maintaining the *moved configuration data while the obtaining configuration information is being performed*.

In rejecting Claim 2, the Examiner acknowledges that neither Kartozi nor Zintel teach such memory configuration. However, the Examiner alleges that Krejsa teaches a system for partitioning memory into different regions and having index corresponding to pluralities of entries in the initialization table. Importantly, the Examiner fails to address the particular details recited in Claim 2 that pertain to the two memory areas - (1) an area for maintaining the previously identified unique identification information, and (2) another area for ***maintaining the moved configuration data while the obtaining configuration information is being performed***. Because none of the references teach or suggest moving configuration data while obtaining configuration information, it logically follows that there is no teaching of a memory area for maintaining such (missing) moved

1443, 1444 (Fed. Cir. 1992). To establish prima facie obviousness of a claimed invention, *all of the claim limitations must be taught or suggested by the prior art*. MPEP 2143.03. *See also, In re Royka*, 490 F.2d 580 (C.C.P.A. 1974). **If the examiner fails to establish a prima facie case, the rejection is improper and will be overturned.** *In re Fine*, 837 F.2d 1071, 1074, 5 USPQ2d

data, as per the features of Claim 2. Again, this moving of data advantageously allows for re-use of discovery code during two distinct discovery operations.

Thus, it is further urged that Claim 2 has been erroneously rejected due to this additional prima facie obviousness deficiency.

D. GROUND OF REJECTION 4 (Claim 21-27)

Claims 21-27 stand rejected under 35 U.S.C. § 103 as being unpatentable over Kartoz in view of Garney (U.S. Patent No. 5,854,905), hereinafter “Garney”.

1. Claim 21

Appellants initially urge error in the rejection of Claim 21 for similar reasons to those given above with respect to Claim 1 (of which Claim 21 depends upon), as the newly cited Garney reference does not overcome the teaching deficiencies identified hereinabove.

Further with respect to Claim 21, such claim recites “wherein the information used to address the particular device comprises a bus identifier that identifies which bus of the plurality of different busses that the particular device is attached to”. Per Claim 1, there is recited a table that comprises (i) an index used to locate particular configuration data for a particular device, (ii) *information used to address the particular device*, and (iii) an offset to a memory location within the particular device at which particular unique identifier information for the particular device is stored. The features of Claim 21 further refine this item (ii) of the table – and in particular, the information used to address a particular device (item (ii) of the table) includes *a bus identifier that identifies which bus of the plurality of different busses that the particular device is attached to*.

In rejecting Claim 21, the Examiner merely alleges:

“Garney teaches a system for managing boot-up and initialization process of a plurality of devices that are connected to multiple different buses”

The Examiner fails to establish – or even allege – that any of the cited references teach/suggest a table that includes information used to address the particular device *comprises a bus identifier that identifies which bus of the plurality of different busses that the particular device is attached to*.

1596, 1598 (Fed. Cir. 1988).

Therefore, Claim 21 has been erroneously rejected due to this prima facie obviousness deficiency. In addition, the burden of proof has not shifted to Appellants to rebut such improper obviousness assertion.⁷

2. *Claim 22*

Appellants initially urge error in the rejection of Claim 22 for similar reasons to those given above with respect to Claim 1 (of which Claim 22 depends upon), as the newly cited Garney reference does not overcome the teaching deficiencies identified hereinabove.

Further with respect to Claim 22, such claim recites “wherein one device of the set of devices contains, in addition to unique identifier information for the one device, identifying information for locating another device of the set of devices within the data processing system”. As can be seen, there is a device that contains both (1) *unique identifier information for such device*, and (2) *identifying information for locating another device*.

In rejecting Claim 22, the Examiner states that all features of Claim 22 are described by Garney at col. 12, lines 36-58. Appellants show that there, Garney states:

“When a bridge has been detected--and initially on a PCI based system, the system BIOS is pre-programmed to detect immediately the PCI bus and initialize it--system BIOS 810 calls for the CPU (not shown) to execute Std_bridge_init 816 to initialize the bridge controller (which in this example, is the PCI controller). Std_bridgeinit module 816 calls the new_ebridge 824 module of PCI BIOS 820 (arrow 1), to initialize the bridge and detect any PCI devices on the system bus. When new.sub_ebridge 824 detects a PCI device with an extensible ROM BIOS A 830, it calls new_edevice 832 module of ROM BIOS A 830 (arrow 2) to identify what kind of device it is. The module new_edevice 832 **returns information about that device** (i.e., is it a bridge device, boot device, etc.) to PCI BIOS 820, and specifically to new_ebridge 824 (arrow 3).

If the device is identified as a bridge device according to the information returned, then new_ebridge module 824 (arrow 5a) calls a ROM BIOS module 830 B of the newly identified bridge device to in turn initialize any devices on that new bridge device. When all the devices on the newly identified bridge have been detected and

⁷ In rejecting claims under 35 U.S.C. Section 103, the examiner bears the initial burden of presenting a prima facie case of obviousness. *In re Oetiker*, 977 F.2d 1443, 1445, 24 USPQ2d 1443, 1444 (Fed. Cir. 1992). Only if that burden is met, does the burden of coming forward with evidence or argument shift to the applicant. *Id.*

initialized, new_ebridge 834 returns program flow control to new_ebridge 824 of PCI BIOS ROM 820 (arrow 4a)”

As can be seen, this cited passage describes initializing a detected bus bridge and the detecting of PCI devices, where a routine (new_edevice) is called to obtain information about a given device (is it a bridge device or boot device, etc). This cited passage does not describe that a given device *itself* contains both (1) *unique identifier information for such device*, and (2) *identifying information for locating another device*.

Therefore, Claim 22 has been erroneously rejected due to this prima facie obviousness deficiency. In addition, the burden of proof has not shifted to Appellants to rebut such improper obviousness assertion.

3. Claim 23

Appellants initially urge error in the rejection of Claim 23 for similar reasons to those given above with respect to Claim 22 (of which Claim 23 depends upon).

Further with respect to Claim 23, such claim recites “wherein the identifying information comprises a bus identifier and an address identifier of where the another device is accessible in the data processing system”. Importantly, Claim 23 depends upon Claim 22, and therefore the features of Claim 23 further refine the two-device identifier aspects of Claim 22, where a given device *itself* contains both (1) *unique identifier information for such device*, and (2) *identifying information for locating another device*. In particular, the features of Claim 23 are directed to further refinements of the (2) *identifying information for locating another device* aspect of Claim 22, where such identifying information for locating the another device (and where such identifying information is contained within a device that also includes unique identifier information for the device itself) includes both (i) a bus identifier and (ii) an address identifier of where such another device is accessible.

The Examiner alleges that all features of Claim 23 are described by Garney at col. 6, line 56 – col. 7, line 14. Appellants show that there, Garney states:

“POST code 500, as shown on FIG. 4, contains several subroutines. First, the CPU is initialized. Next any initialization of the Northbridge takes place. The presence of a video expansion BIOS is detected and the video display controller is initialized.

The memory is tested and initialized next with results displayed to the end user via the video display controller. This sequence of events occurs every time there is a power on, also known as a cold boot, or in some systems also upon a reset, also known as a warm boot. The CPU, memory and Northbridge initializations are grouped together as code block 510 in FIG. 4 since they are consecutively executed. There is also code 520 to initialize the Southbridge, and on PCI systems, code 580 to operate a PCI bridge. The extensible BIOS proposed in the present invention does not need to modify these sections of POST code in order to initialize the system. A keyboard will typically also be detected and its BIOS driver, keyboard driver code 560, initialized before any other expansion BIOS is detected. This allows the expansion BIOS to make use of the keyboard for user input if required. The typical BIOS will also have code that handles and performs boot device support operations, shown in FIG. 4 such as Boot Device ID, selection and policy management code 530. The focus of extensible BIOS is on (1) generic identification and initialization of bridge devices and buses, (2) generic identification and initialization of boot devices, and (3) attaching BIOS device drivers for these generic boot devices to the system BIOS for boot.”

As can be seen, this cited passage describes functions are performed by the POST code. Appellants urge that this description of functions that are performed by POST code does not describe that a given device – which has contained therein a unique identifier that identifies such given device – also includes *identifying information for locating another device*, where such identifying information for locating the another device includes both (i) a bus identifier and (ii) an address identifier of where such *another* device is accessible, as claimed.

Therefore, Claim 23 has been erroneously rejected due to this prima facie obviousness deficiency. In addition, the burden of proof has not shifted to Appellants to rebut such improper obviousness assertion.

4. Claim 24

Appellants initially urge error in the rejection of Claim 24 for similar reasons to those given above with respect to Claim 1 (of which Claim 24 depends upon), as the newly cited Garney reference does not overcome the teaching deficiencies identified hereinabove.

Further with respect to Claim 24, such claim recites “wherein the memory is a volatile memory, and wherein the configuration data is moved from a non-volatile memory to the volatile memory”. As can be seen, the configuration data of Claim 1 is moved from a non-volatile memory to the volatile memory.

In rejecting Claim 24, the Examiner alleges that all features of Claim 24 are taught by Garney at col. 3, lines 44-67. Appellants show that there, Garney states:

“FIG. 2 shows the architecture of a typical personal computer system in accordance with accepted practice. A main system bus 2, such as a Peripheral Component Interconnect (PCI) bus, connects together all the main components of a computer system. A CPU 12 of the computer system is isolated from the system bus 2 by a "Northbridge" 100 which essentially mates the CPU for read and write access to memory 22 (usually RAM). The Northbridge 100 also mates the CPU 12 for read and write access to the system bus 2 for interaction with I/O devices located elsewhere in the system.

The system also has a "Southbridge" 200 which isolates the keyboard 60, the floppy 70 and the hard disk 50 from the system bus 2 and provides a lower cost way to connect these standard devices to the system bus 2 using, for example, an Industry Standard Architecture (ISA) bus. The Southbridge 200 is also used in systems having high performance architectures (such as 32-bit addressing) to free the high performance system bus from dealing directly with add-in devices such as modems or network cards, by connecting to the lower performance (16-bit) ISA bus 4. The lower performance ISA bus 4 connects directly to the modems and network cards that are older in design and cannot take advantage of high performance 32-bit addressing.”

The only ‘operations’ described by the cited passage are (i) a bus that ‘connets together’ all the components, (ii) a CPU is ‘isolated’, (iii) the CPU is ‘mated’ for read and write access to both the memory and the system bus, (iv) a Southbridge that ‘isolates’ peripheral devices from the system bus, and (v) an ISA bus that ‘connects directly’ to communication devices. There is no teaching or suggestion of any type of ‘move’ operation associated with configuration data, and thus this cited passage does not teach/suggest “the configuration data is moved from a non-volatile memory to the volatile memory”, as claimed.

Therefore, Claim 24 has been erroneously rejected due to this prima facie obviousness deficiency. In addition, the burden of proof has not shifted to Appellants to rebut such improper obviousness assertion.

5. Claim 25

Appellants initially urge error in the rejection of Claim 25 for similar reasons to those given above with respect to Claim 24 (of which Claim 25 depends upon).

Further with respect to Claim 25, such claim recites “wherein the current set of configuration data is moved to the non-volatile memory from the volatile memory after being obtained”. As can be seen, per the features of Claim 25, configuration data is obtained, and then after the obtainment of such configuration data, the obtained configuration data is moved to the non-volatile memory from the volatile memory.

In rejecting Claim 25, the Examiner alleges that all features of Claim 25 are described by Garney at col. 4, lines 28-48. Appellants show that there, Garney states:

“For booting, the typical system BIOS will have preinstalled, hardwired code that knows about the Southbridge and configures it to establish the well known boot device interfaces such as keyboard controller, hard disk controller, floppy disk controller. The BIOS checks for the existence of a video expansion BIOS and initializes it if present. A memory check is then performed with the results thereof displayed on the video display. The system BIOS then checks that the hard disk 50 and floppy disk 70 controllers are present at their pre-specified addresses and initializes their drivers. A search for additional expansion BIOS code is then performed. When the network card 250, which connects out to network disk drives, is introduced into a system, the network card 250 will typically be pre-packaged with an expansion BIOS that contains the network boot code. The network boot code is not typically included in the system BIOS. On personal computers of recent vintage, the boot device driver code has been limited to support for a keyboard, a hard disk or floppy disk that is already connected into the PC or that is a permanent feature of the computer.”

As can be seen, this cited passage describes (i) configuring a Southbridge, (ii) checking and initializing video expansion BIOS, (iii) checking memory, (iv) checking if floppy and hard disks are present, and initializing their drivers, and (v) searching for additional expansion BIOS code. This cited passage does not teach or suggest of any type of ‘move’ operation associated with configuration data, and thus this cited passage does not teach/suggest “wherein the current set of configuration data is moved to the non-volatile memory from the volatile memory after being obtained”, as claimed.

Therefore, Claim 25 has been erroneously rejected due to this prima facie obviousness deficiency. In addition, the burden of proof has not shifted to Appellants to rebut such improper obviousness assertion.

6. Claims 26 and 27

Appellants initially urge error in the rejection of Claim 26 (and dependent Claim 27) for similar reasons to those given above with respect to Claim 1 (of which Claim 26 depends upon), as the newly cited Garney reference does not overcome the teaching deficiencies identified hereinabove.

Further with respect to Claim 26 (and dependent Claim 27), such claim recites “wherein the identifying step is performed by an embedded processor of the data processing system while a plurality of primary processors of the data processing system are powered-off”. As can be seen, the features of Claim 26 are directed to a further refinement of the ‘identifying’ step of Claim 1, where unique identification information for a set of devices in the data processing system are identified. Per Claim 26, the identification of such unique identification information for the devices is performed by an *embedded processor of the data processing system while a plurality of primary processors of the data processing system are powered-off*.

In rejecting Claim 26, the Examiner alleges that all features of Claim 26 are described by Garney at col. 6, lines 43-55. Appellants show that there, Garney states:

“FIG. 4 illustrates the code contained within a typical BIOS. In a typical BIOS, there are a set of routines that handle events after reset and power-on of the computer but before the invocation of an operating system (such as Disk Operating System (DOS)). This set of routines that boot the computer system prior to loading the operating system is known as Power-On-Self-Test (POST). Certain devices within the computer system must always be started and in a particular order to maintain functionality and accessibility throughout the boot-up process. For instance prior to even testing the memory (RAM), there is a need to invoke the video BIOS routines so that when the memory test does run, the user can monitor its progress through a display device.”

As can be seen, this cited Garney description describes BIOS code routines, including (i) handling events after reset and power-on of the computer, and (ii) the starting of internal devices in a particular order – such as invoking a video BIOS routine prior to testing the memory. This cited

passage does not describe any type of unique identification information being identified, and therefore cannot describe or teach the particular characteristics of such an ‘identifying’ step, such as performing such unique identification information *by an embedded processor of the data processing system while a **plurality** of primary processors of the data processing system are powered-off*, as claimed. Importantly, Garney’s Figure 2 depicts a single CPU 12.

Therefore, Claim 26 (and similarly for Claim 27) has been erroneously rejected due to this prima facie obviousness deficiency. In addition, the burden of proof has not shifted to Appellants to rebut such improper obviousness assertion.

E. CONCLUSION

As shown above, the Examiner has failed to state valid rejections against any of the claims. Therefore, Appellants request that the Board of Patent Appeals and Interferences reverse the rejections. Additionally, Appellants request that the Board direct the examiner to allow the claims.

Date: March 9, 2009

Respectfully submitted,

/Wayne P. Bailey/

Wayne P. Bailey
Reg. No. 34,289
Yee & Associates, P.C.
P.O. Box 802333
Dallas, TX 75380
(972) 385-8777

CLAIMS APPENDIX

The text of the claims involved in the appeal is as follows:

1. A method in a data processing system for identifying device configurations, the method comprising:

identifying unique identification information for a set of devices in the data processing system to form identified unique identification information;

comparing the identified unique identification information with previously identified unique identification information;

moving configuration data to a memory for devices in the set of devices in which a match exists between the identified unique identification information and the previously identified unique identification information for devices; and

obtaining configuration information from a device in which configuration information is absent in the memory after configuration data has been moved to the memory for the devices to form a current set of configuration data for the set of devices, wherein the previously identified unique identification information is accessed using a table associated with the configuration data for the set of devices, wherein the table comprises (i) an index used to locate particular configuration data for a particular device, (ii) information used to address the particular device, and (iii) an offset to a memory location within the particular device at which particular unique identifier information for the particular device is stored.

2. The method of claim 1, wherein the memory is a temporary random access memory comprising an area for maintaining the previously identified unique identification information and another area for maintaining the moved configuration data while the obtaining configuration information is being performed.
3. The method of claim 1, wherein the unique identification information is a unique device identifier.
4. The method of claim 1, wherein the current configuration data for the set of devices is stored in a set of files.
5. The method of claim 1, wherein the unique identification information is identified by reading the unique identification information from the set of devices.
21. The method of Claim 1, wherein the data processing system comprises a plurality of different busses, and wherein the information used to address the particular device comprises a bus identifier that identifies which bus of the plurality of different busses that the particular device is attached to.
22. The method of Claim 1, wherein one device of the set of devices contains, in addition to unique identifier information for the one device, identifying information for locating another device of the set of devices within the data processing system.

23. The method of Claim 22, wherein the data processing system comprises a plurality of different busses, and wherein the identifying information comprises a bus identifier and an address identifier of where the another device is accessible in the data processing system.

24. The method of Claim 1, wherein the memory is a volatile memory, and wherein the configuration data is moved from a non-volatile memory to the volatile memory.

25. The method of Claim 24, wherein the current set of configuration data is moved to the non-volatile memory from the volatile memory after being obtained.

26. The method of Claim 1, wherein the identifying step is performed by an embedded processor of the data processing system while a plurality of primary processors of the data processing system are powered-off.

27. The method of Claim 26, wherein the obtaining step is initiated by the embedded processor during an initial program load (IPL) of the data processing system.

EVIDENCE APPENDIX

This appeal brief presents no additional evidence.

RELATED PROCEEDINGS APPENDIX

This appeal has no related proceedings.